



(19) **United States**

(12) **Patent Application Publication**
Pettersson et al.

(10) **Pub. No.: US 2014/0294362 A1**

(43) **Pub. Date: Oct. 2, 2014**

(54) **METHOD AND APPARATUS FOR
DETECTING QUALITY DEFECTS IN A
VIDEO BITSTREAM**

Publication Classification

(75) Inventors: **Martin Pettersson**, Vallentuna (SE);
Savvas Argyropoulos, Bonn (DE)

(51) **Int. Cl.**
H04N 7/64 (2006.01)
(52) **U.S. Cl.**
CPC **H04N 7/64** (2013.01)
USPC **386/264**

(73) Assignees: **Telefonaktiebolaget L M Ericsson
(PUBL)**, Stockholm (SE); **Deutsche
Telekom AB**, Bonn (DE)

(57) **ABSTRACT**

(21) Appl. No.: **14/363,607**

There is provided an apparatus for detecting a quality defect in a video sequence, the apparatus comprising: a receiver arranged to receive a video bitstream representing a video sequence; and a defect module arranged to determine a measure of quantization parameter for each picture of the video sequence, the defect module further arranged to identify a quality defect when an abrupt change in the measure of quantization parameter occurs.

(22) PCT Filed: **Dec. 9, 2011**

(86) PCT No.: **PCT/EP2011/072288**

§ 371 (c)(1),
(2), (4) Date: **Jun. 6, 2014**




Fig. 1c




Fig. 1b




Fig. 1a




Fig. 2

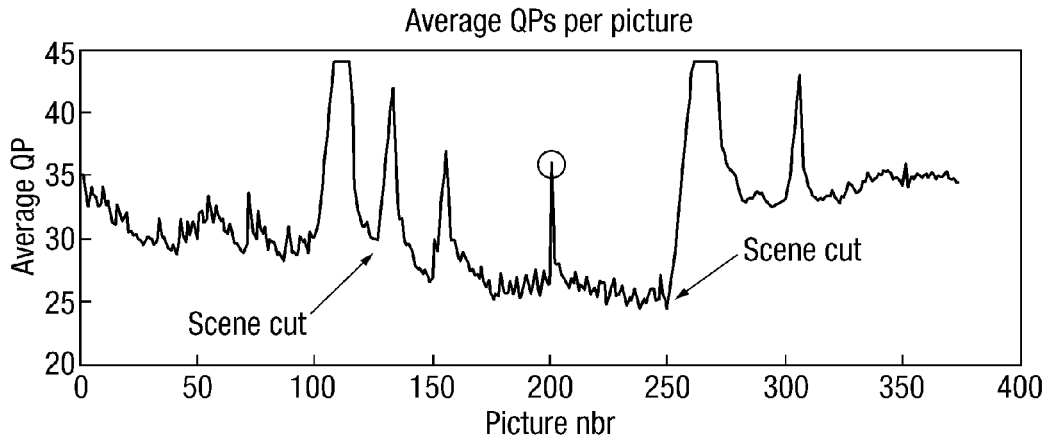


Fig. 3




Fig. 4a




Fig. 4b




Fig. 4c




Fig. 4d




Fig. 5




Fig. 6

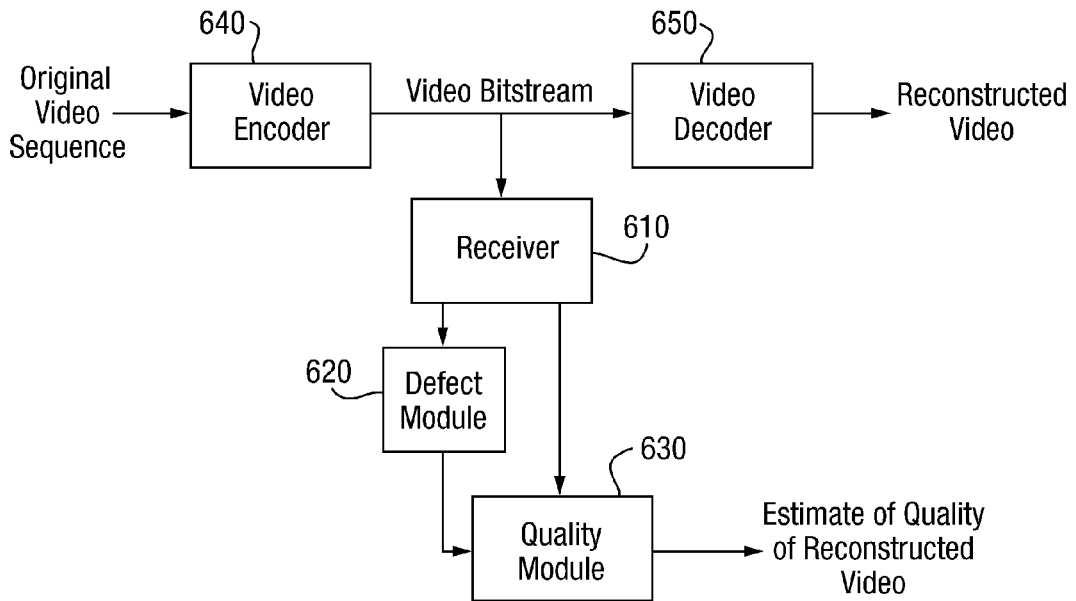


Fig. 7





Fig. 8



METHOD AND APPARATUS FOR DETECTING QUALITY DEFECTS IN A VIDEO BITSTREAM

TECHNICAL FIELD

[0001] The present application relates to a method for detecting a quality defect in a video sequence; an apparatus for detecting a quality defect in a video sequence; and a computer-readable medium.

BACKGROUND

[0002] Block-based coding is the dominant video encoding technology with codec standards such as H.263, MPEG-4 Visual, MPEG-4 AVC (a.k.a. H.264) and the emerging HEVC/H.265 standard being developed in JCT-VC. These codecs use different types of pictures (which employ different types of prediction) to compress the video as efficiently as possible. An intra-coded picture (I-picture) may only be predicted spatially from areas within the picture itself. A predictive picture (P-picture) is predicted from previously coded I or P pictures. A bidirectional predictive pictures (B-picture) is predicted from both previous and/or subsequent pictures.

[0003] The result of predicting a particular picture from other pictures is likely to produce a result which is different to the original version of the particular picture. The differences between the prediction of a picture and the original picture are called the picture's residual data. The residual data is transmitted or stored together with the prediction instructions in order to allow the original picture to be accurately reconstructed. In order to make the residual data more compact the data is transformed into the frequency domain. In H.264/AVC the Hadamard transform and a transform similar to a Discrete Cosine Transform are used on 4×4 or 2×2 (chroma) blocks for this purpose. Once the residuals have been transformed a quantization is performed on the transformed data to limit the data that needs to be transmitted or stored. It is in this step that the actual lossy compression takes place. The level of quantization is determined by the quantization parameter (QP) and a corresponding look-up-table. The QP is set at picture level but can also be altered at macro-block (MB) level.

[0004] A higher QP means that the residual data are more coarsely quantized, so less detail of the residual data is captured and so the reconstructed picture will match the original picture less well. A lower QP means that more detail of the residual data is captured and so the reconstructed picture will match the original picture well.

[0005] It is fairly expensive in terms of bits for an encoder to alter the QP at MB level and so the majority of changes in QP occur between pictures. When the quantization has been performed the quantized transformed residuals are encoded using entropy coding. In H.264/AVC one of the entropy coding algorithms, namely CAVLC (Context Adaptive Variable Length Coding) or CABAC (Context Adaptive Binary Arithmetic Coding), is used for this.

[0006] To increase error resilience in error prone networks I-pictures are usually inserted periodically to refresh the video. I-pictures are also inserted periodically to allow for random access and channel switching. Where the forced intra pictures should be inserted in time is not defined by the video coding standard, but is up to the encoder to decide. Typically, video coding standards define the video bitstream syntax and the decoding process, but do not define the encoding process.

In other words, the method by which the video sequence is encoded is not standardized, whereas the output of the encoding process is.

[0007] To ensure the end-to-end quality of video over fixed and mobile networks network operators and broadcast vendors can utilize objective video quality models. Objective video quality models are mathematical models that approximate results of subjective quality assessment, but are based on criteria and metrics that can be measured objectively and automatically evaluated by a computer program.

[0008] The performance of an objective video quality model is evaluated by computing a metric between the objective score generated by the objective video quality model and subjective test results. This metric can be, for example, the correlation between subjective and objective data or the mean squared error. A subjective test result may comprise a mean opinion score (MOS) obtained from the opinions of a plurality of human test subjects.

[0009] Perceptual models may be considered to be a subset of objective video quality models. Whereas an objective video quality model can refer to any automated quality assessment method, a perceptual model attempts to determine to what extent any quality defects would be perceived by a viewer. Perceptual models can utilize the pixel information in the decoded video sequence, and in the case of full-reference models the reference signal may also be used to predict the degradation of the processed video. A big disadvantage of perceptual models is that they are usually computational demanding and not very suitable for deployment on a large scale for network monitoring purposes.

[0010] An alternative approach to perceptual quality models that is more light-weight than a full-reference model is to use network layer protocol headers as input for quality estimation of a transmitted video. This approach makes the model very efficient to implement and use, but the quality estimation of the transmitted video will be rather coarse. Therefore a video bitstream quality model may also be implemented. This model takes the encoded elementary stream as input in addition to network protocol headers and has the advantage that it will be fairly light-weight and yet has the potential of getting a better estimate of the quality of the video than one just using network layer protocol headers. Such a video bitstream quality model may operate in two modes, one mode where full decoding of the bitstream is allowed and another, lower complexity, mode where full decoding is not allowed (such that pixel information cannot be used).

[0011] "Modeling the impact of frame rate and quantization stepsizes and their temporal variations on perceptual video quality: a review of recent works" by Y-F. Qu, Z. Ma, and Y. Wang, and published in Information Sciences and Systems (CISS), 2010 44th Annual Conference, 17-19 Mar. 2010, describes the effect of frame rate and quantization stepsize as well as the temporal variation of the frame rate on the perceptual quality.

[0012] "Evaluation of temporal variation of video quality in packet loss networks" by C. Yim and A. Bovik, and published in Signal Processing: Image Communication 26 (2011) describes the effect that variations in the temporal quality of videos have on global video quality.

[0013] There is thus a need for quality defect detection that does not require decoding of the encoded video bitstream. Such a quality defect detection may be suitable for implementation with a quality model.

SUMMARY

[0014] It has been recognized that abrupt changes in the Quantization Parameter (QP) of a block-coded video stream can give rise to notable quality defects in a decoded video sequence. A method and apparatus for detecting such defects is described which does not require the decoding of the encoded video sequence from the video bitstream. The estimation of the visual impact of the change in QP can be used in conjunction with other model parameters in a video bitstream quality model to estimate the overall quality of a video bitstream. The solution described herein thus improves the accuracy of the estimation. This allows an improved video bitstream quality model without the computational overhead of decoding the video sequence from the video bitstream.

[0015] Accordingly, there is provided a method for detecting a quality defect in a video sequence. The method comprises receiving a video bitstream representing a video sequence. The method further comprises determining a measure of quantization parameter for each picture of the video sequence. The method further comprises identifying a quality defect when an abrupt change in quantization parameter occurs.

[0016] Certain quality defects can be identified in a video bitstream representing a video sequence without decoding the video sequence from the video bitstream. This means quality defects can be identified in the video bitstream at network locations without a decoder, which allows for more accurate quality estimation at nodes throughout the distribution network.

[0017] The method may further comprise determining if each picture of the video sequence is an intra-coded picture, and identifying a quality defect as intra-picture flicker when an abrupt change in quantization parameter occurs for an intra-coded picture.

[0018] Intra-picture flicker is one type of quality defect that is caused by an abrupt change in quantization parameter. Intra-picture flicker may be identified when an intra coded picture coincides with a peak in quantization parameter. Intra-picture flicker may be identified when an abrupt change in quantization parameter occurs between an intra-coded picture and its neighbouring inter-coded pictures.

[0019] The method may be used in a quality estimation model to determine a perceptual quality of the video sequence represented by the video bitstream.

[0020] There is further provided an apparatus for detecting a quality defect in a video sequence, the apparatus comprising a receiver and a defect module. The receiver is arranged to receive a video bitstream representing a video sequence. The defect module is arranged to determine a measure of quantization parameter for each picture of the video sequence, and is further arranged to identify a quality defect when an abrupt change in the measure of quantization parameter occurs.

[0021] The apparatus may further comprise a picture type determining module arranged to determine whether each picture of the video sequence is an intra-coded picture. The defect module may be further arranged to identify a quality defect as an intra-picture flicker when an abrupt change in the measure of quantization parameter occurs for an intra-coded picture.

[0022] The apparatus may further comprise a quality module arranged to receive the video bitstream and a defect indication from the defect module, the quality module further arranged to apply a quality model to estimate a quality of video reconstructed from the video bitstream.

[0023] There is further provided a computer-readable medium, carrying instructions, which, when executed by computer logic, causes said computer logic to carry out any of the methods defined herein.

BRIEF DESCRIPTION OF THE DRAWINGS

[0024] A method and apparatus for detecting a quality defect in a video bitstream will now be described, by way of example only, with reference to the accompanying drawings, in which:

[0025] FIG. 1 shows sample sections of three pictures from one scene within a video sequence;

[0026] FIG. 2 shows a plot of the average QP per picture for the video sequence of FIG. 1;

[0027] FIG. 3 shows a plot for the average QP per picture for a second video sequence;

[0028] FIG. 4 illustrates four types of abrupt change in QP; **[0029]** FIG. 5 illustrates an apparatus for performing the method described herein;

[0030] FIG. 6 illustrates an apparatus incorporating the detection of an abrupt change in QP into a quality model;

[0031] FIG. 7 illustrates a method described herein; and

[0032] FIG. 8 illustrates another method described herein.

DETAILED DESCRIPTION

[0033] One example of a quality defect caused by an abrupt change in quantization parameter (QP) creating a perceptible quality defect is intra-picture flicker. Intra-picture flicker is a coding artifact that is usually due to the bit rate control of the encoder not being able to keep its requirements of having an almost constant quality given a certain bitrate. Bit rate control in video coding is a complex task and is not a part of any standard. It is up to the developer of the encoder to create a rate control that allocates the available bit budget optimally over the pictures of the sequence to maximize the perceived quality. Rate control requirements can be very different. Rate controls for video-calling are often strict to keep the delay of the video low. This sometimes means keeping to a certain bit budget for each picture. One-way video, like broadcasted mobile TV, usually has more relaxed bit rate control requirements with a bit budget for a certain time period, e.g. 1-2 seconds. More flexibility in the rate of encoded video can be provided by implementing statistical multiplexing and distributing a fixed amount of bandwidth across a plurality of parallel video sequences in response to the encoding demands of each sequence. In some encoding arrangements the encoder does not have access to the subsequent pictures and it must select the encoding parameters on-the-fly (which is also known as one-pass encoding). For these and other reasons, during encoding the bitrate may not be optimally allocated between the pictures of the video sequence.

[0034] Intra-picture flicker arises when periodic intra-pictures are used when encoding a scene. The artifact is mainly visible at medium to low bitrates, where the requirements of the rate control can be more difficult to meet. However, the artifact may also be visible even if high bitrates are used. Intra-picture flicker is also more apparent for static scenes and where the inter (P- and B-) pictures are relatively inexpensive to encode in terms of bits (but the intra-pictures still are expensive) and the artifact is not masked by motion within the scene.

[0035] A typical intra-picture flicker artifact is illustrated by FIG. 1. FIG. 1 shows magnified sections of three pictures

from one scene within a video sequence, A. FIG. 1a is a portion of picture 199 of sequence A, this picture is an inter-coded picture encoded with an average QP of 26. FIG. 1b is a portion of picture 200 of sequence A, this picture is an intra-coded picture encoded with an average QP of 36. FIG. 1c is a portion of picture 201 of sequence A, this picture is an inter-coded picture encoded with an average QP of 28. The higher QP in picture 200 of the video sequence is apparent from the blurriness of this picture compared to the other two pictures.

[0036] The content of video sequence A is of a relatively static nature and depicts grazing cattle. The sequence has been encoded and decoded using JM, which is the reference encoder for H.264/AVC. The pictures 199 and 201 have been encoded as inter-pictures with a relatively low QP. The middle picture 200 has been encoded as an intra-picture with a relatively high QP due to the restrictions of the rate control and the fact that intra-pictures are much more expensive to encode than inter-pictures. Looking at the pictures played out in a sequence will show an obvious and annoying flicker, where the spatial detail within the image disappears at picture 200, but then re-appears at picture 201.

[0037] Such intra-picture flicker can be seen in video bitstreams from a variety of encoders. It is not a part of the encoder bit-rate control strategy, but rather a failure to follow the strategy. Since video codec standards do not cover the behaviour of the encoder, but rather the syntax of the bitstream to be decoded by the decoder, there are many different encoding solutions out there. Some encoding solutions manage the intra-picture flickering problem better than others.

[0038] Detecting such a defect is particularly important when statistical multiplexing is used and an easy-to-encode channel does not get enough bits to handle the periodic intra-frames correctly. One could also expect the problem to look worse at lower resolutions (e.g. the mobile TV case), when the available bitrate is less, and the encoding complexity per frame has less variation. For instance, having a periodic intra period of 0.5 seconds at a “medium” bitrate for HDTV is usually ok, while using a periodic intra period of 0.5 seconds at a “medium” bitrate for mobile TV (e.g. 320x240 pixels) would very likely result in regular intra-picture flickering.

[0039] Presented herein is a method and apparatus for detecting when this quality dip appears by monitoring the QP value. If the QP value is increased by a certain amount between an inter-picture and a following intra-picture and decreased a certain amount for the next-coming inter picture, then an intra-picture flicker is detected. In a preferred embodiment the QP is averaged over each picture before the comparison is made. In another embodiment, the QP difference is calculated per macro block and then the difference is averaged over all macro-blocks.

[0040] FIG. 2 shows a plot of the average QP per picture for sequence A. An intra-picture flicker is present at picture number 200. The peak in QP which gives rise to the intra-picture flicker is marked with a circle on the plot. The sequence contains two scene cuts, marked by arrows, at picture numbers 130 and 250. Despite the fact that the QP is increased a lot when a scene cut occurs this is not annoying to the user and so is not interpreted as an obvious quality defect. This is because the scene is changed and the quality difference between pictures from the old scene and the new scene will not be as apparent to the viewer. Still, the case where the encoded picture quality suddenly drops (either after a scene cut or in the middle of a scene) can be identified by detecting

when the QP is increased between an inter-picture and an intra-picture without decreasing considerably after that.

[0041] The plot of FIG. 2 also shows a few more peaks at positions ~110, ~155 and ~305. Since the QP is increasing over a few inter-pictures here this is hardly visible to a viewer. Inter-pictures are predicted from previous pictures and so a decrease in fidelity of the residuals for a few pictures will not do much to the apparent quality (while the spatial detail may not be accurate, an appropriate amount of spatial detail is expected to be imported from the referred to pictures). The pseudo code below shows how intra-picture flicker and corresponding sudden drops in encoding quality can be detected:

```

if (picn == inter and picn+1 == intra and QPn+1 - QPn >
threshold t1) {
  if (picn+1 == scene_cut)
    drop in encoding quality in new scene detected
  else if (QPn+1 - QPn+2 > threshold t2) {
    intra picture flicker detected
  } else {
    drop in encoding quality detected
  }
}

```

[0042] FIG. 3 shows a plot for the average QP per picture for a second video sequence, B. This sequence depicts a newsreader where mostly just his lips are moving. Every 50th picture is encoded as an intra-coded picture. Since the scene is very static, the inter-predictions become relatively accurate with a low amount of residual data and a relatively low QP can be afforded for the inter-pictures. The encoder can thus encode the inter-pictures using a relatively low bit rate. However, when a periodic intra-picture is forced into the encoding process, it is much more expensive to encode in terms of bits and the QP must be raised drastically to accommodate it without causing a sharp spike in bitrate demand. Furthermore, inserting the periodic intra-pictures also sets the rate control in imbalance resulting in an increase in QP for the inter-pictures following after a periodic intra picture. Five large intra picture flickers are found in this sequence and are marked with circles in the plot of FIG. 3.

[0043] The high QP peaks a few pictures after the periodic intra pictures do not affect quality that much for the same reason described for sequence A: inter-pictures are predicted from previous (good quality) pictures and a decrease in fidelity of the residuals will not do much to the perceived quality over a few inter-pictures.

[0044] The range in QP over a sequence also affects the distribution of the quality of a sequence. The span in quality can be calculated as:

$$QP_{range} = QP_{max} - QP_{min}$$

[0045] Intra-picture flicker is just one example of an abrupt change in QP causing a quality defect in an encoded video sequence. FIG. 4 illustrates four types of abrupt change in QP: a peak in FIG. 4a; a trough in FIG. 4b; a step up in FIG. 4c; and a step down in FIG. 4d. A peak is equivalent to a step up immediately followed by a step down, and a trough is equivalent to a step down immediately followed by a step up.

[0046] Intra-picture flicker as described above occurs when a peak as shown in FIG. 4a coincides with an intra-coded picture. Such a peak is likely to coincide with an intra-coded picture for the reasons given above, but such a peak could occur coincident with an inter-coded picture, and if it did it is

likely that the sharp decrease in spatial detail of the scene would be apparent to the viewer.

[0047] A trough as shown in FIG. 4*b* will be apparent to a viewer as a sudden increase in spatial detail within a scene, which could be as jarring to the viewer as an intra-picture flicker. Such a trough would be apparent whether or not it coincided with an intra-picture.

[0048] Even a step up as illustrated in FIG. 4*c* would be apparent to a viewer as a sudden decrease in resolution of spatial detail (a sudden increase in blurriness) in the scene. Such a step-up in QP is more likely to be noticeable if it coincides with an intra-picture.

[0049] Similarly, a step down as illustrated in FIG. 4*d* would be apparent to a viewer as a sudden increase in resolution of spatial detail (a sudden decrease in blurriness) in the scene. Such a step-down in QP is more likely to be noticeable if it coincides with an intra-picture.

[0050] It has been shown how an intra-picture flicker is detected within a video bitstream without fully decoding the video bitstream. It has also been shown how the general case of a sudden change in QP is detected. This information can be used to improve the accuracy of the quality estimation from a video bitstream quality model.

[0051] FIG. 5 illustrates an apparatus for performing the method described herein. The apparatus comprises a receiver 510, and a defect module 520. The receiver 510 is arranged to receive a video bitstream representing a video sequence. The defect module 520 is arranged to determine a measure of an average quantization parameter for each picture of the video sequence, and to identify a quality defect when an abrupt change in the measure of quantization parameter occurs. When a quality defect is identified by the defect module 520, it outputs a defect detection indication.

[0052] FIG. 6 illustrates an apparatus incorporating the detection of an abrupt change in QP into a quality model. The apparatus comprises a receiver 610, a defect module 620, and a quality module 630. The receiver 610 is arranged to receive a video bitstream representing a video sequence. The video bitstream is generated by a video encoder 640 by encoding an original video sequence. The video bitstream is also received by a video decoder 650 where the video bitstream is decoded and a reconstructed video sequence is generated. The defect module 620 determines a measure of an average quantization parameter for each picture of the video sequence, and identifies a quality defect when an abrupt change in the measure of quantization parameter occurs. The quality module 630 receives the video bitstream and a defect indication from the defect module 620. The quality module 630 uses these inputs to apply a quality model and to estimate a quality of the video reconstructed from the video bitstream.

[0053] A video bitstream quality model may further take into account one or more of the following:

[0054] Spatial detail in the scene

[0055] Amount of movement in the scene

[0056] Existence of intra picture flicker

[0057] Number of intra picture flicker occasions

[0058] Frequency of intra picture flicker (per picture or per second)

[0059] Distance (in terms of time or number of pictures) between intra picture flicker events

[0060] Difference in QP for intra picture flicker

[0061] Difference in QP for a drop in encoding quality starting at an intra-picture

[0062] Duration of drop in encoding quality starting at an intra-picture

[0063] Difference in QP for a raise in encoding quality starting at an intra-picture

[0064] Duration of raise in encoding quality starting at an intra-picture

[0065] Difference in QP for a QP peak over inter pictures

[0066] Range of QP for the sequence

[0067] One or more of these details may be included in the calculation of the quality estimation of a video bitstream. How the total quality model is constructed is not in the scope of this document, but it could for instance be a linear model like an additive model or a non-linear model like a multiplicative model.

[0068] For example, the spatial detail in the picture may be determined using frequency analysis such as Fourier analysis or a discrete cosine transform. The magnitude and number of high frequency components gives a measure of the spatial detail. Since block-based video codecs transform image and residual data into the frequency domain, it is possible to obtain a measure of the spatial detail of a picture by parsing out from the video bitstream the coefficients relating to the frequency components. Such coefficients are often called AC coefficients. The measure of the spatial detail can be used by a quality model to determine the impact of the quality defect on the quality score. Where the measure of spatial detail is high then the quality defect will have a greater impact on the quality score.

[0069] As a further example, a measure of the amount of movement in the picture may be determined by the average magnitude of the motion vectors for macroblocks in the picture. The measure of the amount of movement may be used by a quality model to determine the impact of the quality defect on the quality score. Where the measure of the amount of movement is high then the quality defect will have a lower impact on the quality score.

[0070] FIG. 7 illustrates a method described herein. At 710 a video bitstream representing a video sequence is received. At 720, a measure of an average quantization parameter for each picture of the video sequence is determined. At 730, a quality defect is identified when an abrupt change in quantization parameter occurs. This quality defect may be used by a quality model to estimate the quality of video reconstructed from the video bitstream.

[0071] FIG. 8 illustrates another method described herein. At 810 a video bitstream representing a video sequence is received. At 820, a measure of an average quantization parameter for each picture of the video sequence is determined. At 830, it is determined whether each picture is an I-picture. At 840, an intra-picture flicker is identified when the current picture is identified as an I-picture and the quantization parameter for the current picture is both a threshold amount greater than the quantization parameter of the preceding, and the subsequent, picture. The identification of such intra-picture flickering may be used by a quality model to estimate the quality of video reconstructed from the video bitstream.

[0072] An advantage of the method and apparatus described herein is that these allow detection of intra-picture flicker and similar quality defects without decoding the bitstream. This means that sudden drops in encoding quality are detected. The estimate of the visual impact of these quality defects can be used to improve the accuracy of the estimated quality from a video bitstream quality model. Further,

because the video bitstream quality model does not need to decode the bitstream to determine the quality, video quality monitoring can be performed at parts of the distribution network where decoding the video bitstream would be too costly, in terms of processing resources, power consumption, physical space, etc.

[0073] It will be apparent to the skilled person that the exact order and content of the actions carried out in the method described herein may be altered according to the requirements of a particular set of execution parameters. Accordingly, the order in which actions are described and/or claimed is not to be construed as a strict limitation on order in which actions are to be performed.

[0074] Further, while examples have been given in the context of particular communications standards, these examples are not intended to be the limit of the communications standards to which the disclosed method and apparatus may be applied. For example, while specific examples have been given in the context of H.264, the principles disclosed herein can also be applied to MPEG-2, H.263, MPEG-4, HEVC/H.265 or other codec, and indeed any video codec using quantization of transform coefficients.

1. A method for detecting a quality defect in a video sequence, the method comprising:

receiving a video bitstream representing a video sequence; determining a measure of quantization parameter for each picture of the video sequence; and identifying a quality defect when an abrupt change in quantization parameter occurs.

2. The method of claim 1, further comprising:

determining if each picture of the video sequence is an intra-coded picture; and identifying a quality defect as intra-picture flicker when an abrupt change in quantization parameter occurs for an intra-coded picture.

3. The method of claim 1, wherein the abrupt change in quantization parameter comprises a measure of quantization parameter for a picture that is greater than, by a threshold amount, the measure of quantization parameter of a predetermined number of preceding pictures.

4. The method of claim 1, wherein the abrupt change in quantization parameter comprises a measure of quantization parameter for a picture that is less than, by a threshold amount, the measure of quantization parameter of a predetermined number of preceding pictures.

5. The method of claim 1, wherein the abrupt change in quantization parameter comprises a measure of quantization parameter for a picture that is greater than, by a threshold amount, the measure of quantization parameter of a first predetermined number of preceding pictures and a second predetermined number of subsequent pictures.

6. The method of claim 1, wherein the abrupt change in quantization parameter comprises a measure of quantization parameter for a picture that is less than, by a threshold amount, the measure of quantization parameter of a first

predetermined number of preceding pictures and a second predetermined number of subsequent pictures.

7. The method of claim 5, wherein the first predetermined number of pictures is equal to the second predetermined number of pictures.

8. The method of claim 3, wherein the predetermined number of pictures is one.

9. The method of claim 1, further comprising:

determining the impact of the quality defect by further measuring at least one of:

the difference in the measure of quantization parameter; the spatial detail in the picture; and the amount of movement in the picture.

10. The method of claim 1, wherein the measure of quantization parameter for a picture is the average quantization parameter used for the picture.

11. The method of claim 1, wherein the method is used in a quality estimation model to determine a perceptual quality of the video sequence represented by the video bitstream.

12. The method of claim 1, wherein the video bitstream is encoded using at least one of: H.264, MPEG-2, H.263, MPEG-4 and HEVC/H.265.

13. An apparatus for detecting a quality defect in a video sequence, the apparatus comprising:

a receiver arranged to receive a video bitstream representing a video sequence; and

a defect module arranged to determine a measure of quantization parameter for each picture of the video sequence, the defect module further arranged to identify a quality defect when an abrupt change in the measure of quantization parameter occurs.

14. The apparatus of claim 13, further comprising:

a picture type determining module arranged to determine whether each picture of the video sequence is an intra-coded picture; and

wherein the defect module is further arranged to identify a quality defect as an intra-picture flicker when an abrupt change in the measure of quantization parameter occurs for an intra-coded picture.

15. The apparatus of claim 13, further comprising a quality module arranged to receive the video bitstream and a defect indication from the defect module, the quality module further arranged to apply a quality model to estimate a quality of video reconstructed from the video bitstream.

16. The apparatus of claim 15, wherein the quality module is further arranged to determine the magnitude of the quality defect by further measuring at least one of:

the difference in the measure of quantization parameter; the spatial detail in the picture; and the amount of movement in the picture.

17. A non-transitory computer-readable medium, carrying instructions, which, when executed by computer logic, causes said computer logic to carry out the method of claim 1.

* * * * *